

CSL overview for ACEs

Ben Kilminster, Kevin McFarland, Tony Vaiciulis
University of Rochester

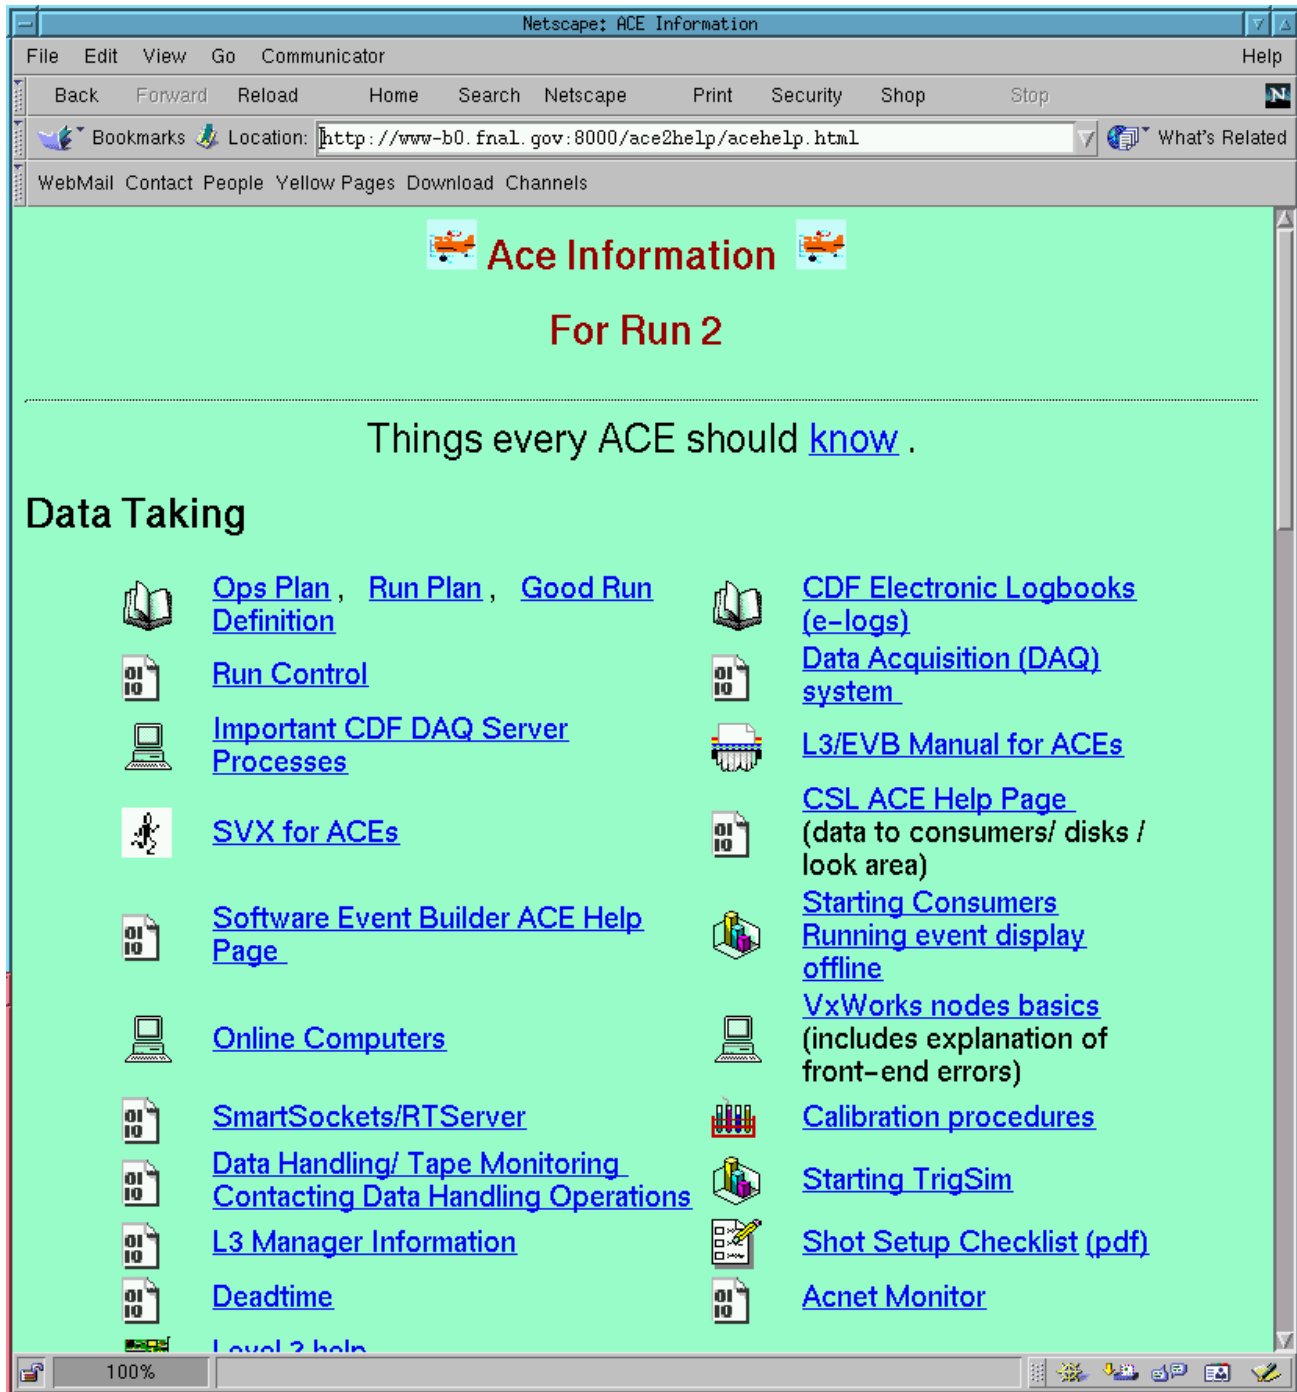
Hirokazu Kobayashi, Hiroyuki Matsunaga, Makoto Shimojima
University of Tsukuba

January 16, 2002

- CSL description
- monitoring
- troubleshooting

CSL ACE Help Page

Most items in this talk are documented on CSL ACE help page



CSL ACE Help Page

Netscape: Consumer-Server/Logger ACE Help Page

File Edit View Go Communicator Help

Back Forward Reload Home Search Netscape Print Security Shop Stop

Bookmarks Location: <http://www-b0.fnal.gov:8000/ace2help/csl/> What's Related

WebMail Contact People Yellow Pages Download Channels

Consumer-Server/Logger ACE Help Page

About the CSL

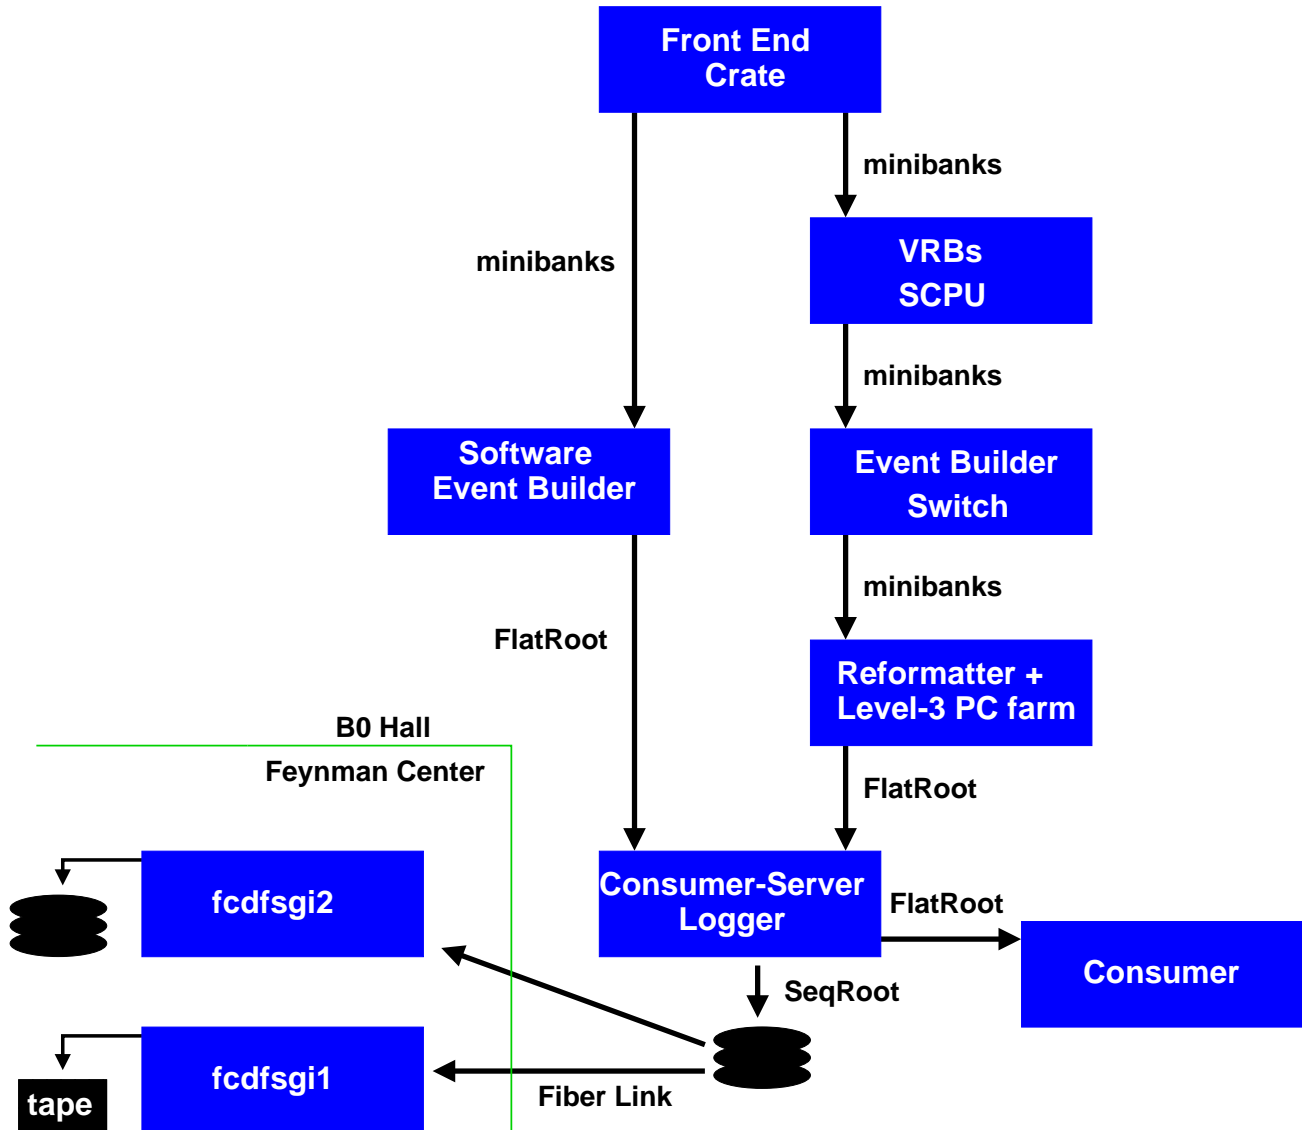
Questions, comments, suggestions? Send mail to cdf_csl@fnal.gov

The Consumer-Server/Logger (CSL) is the central online hub through which all CDF data taken during Run II must pass. Events passing the Level-3 trigger are written to disk and served to consumers by the CSL. The raw data files on disk are then moved to the Feynman Computing Center (FCC) for further processing. Consumers are online monitoring processes, spying on a fraction of the data passing through the CSL. The diagram below shows how the CSL fits into the overall dataflow.

```
graph TD;
    FEC[Front End Crate] -- minibanks --> SEB[Software Event Builder];
    FEC -- minibanks --> VRB[VRBs SCPU];
    VRB -- minibanks --> EBS[Event Builder Switch];
    EBS -- minibanks --> RLF[Reformatter + Level-3 PC farm];
    SEB -- FlatRoot --> RLF;
    RLF --> B0[B0 Hall];
    B0 --> FCC[Feynman Center];
```

The diagram illustrates the dataflow of the Consumer-Server/Logger (CSL). It shows the path from the Front End Crate through various processing stages to the B0 Hall and Feynman Center. The Front End Crate sends data via minibanks to both the Software Event Builder and the VRBs SCPU. The VRBs SCPU sends data via minibanks to the Event Builder Switch, which then sends data via minibanks to the Reformatter + Level-3 PC farm. The Software Event Builder sends data via FlatRoot to the Reformatter + Level-3 PC farm. The Reformatter + Level-3 PC farm sends data to the B0 Hall, which then sends data to the Feynman Center.

CSL description



primary CSL functions

- receive events from Level-3 PC farm at 20 MB/sec ($75 \text{ Hz} \times 250 \text{ kB/event}$)
- receive events from the software event builder
- write events to disk at $\sim 20 \text{ MB/sec}$
- handle as many consumer requests as possible (5-10 MB/s total)

Starting/stopping the CSL

During normal running, the CSL never needs to be started or stopped. It is always "on", ready to receive events.

The ace is able to start or stop the CSL processes.

The following commands are available:

b0dau32: cslcom stop (stop all CSL processes)

b0dau32: cslcom check (check if CSL processes are running)

b0dau32: cslcom start (start all CSL processes)


b0dau32: cslcom cleanup (kill all CSL processes, normally you don't need to do this)

Whenever start, stop, or cleanup are executed mail is automatically sent to the CSL group. Shift crew should also send a mail to

`cdf_csl@fnal.gov`

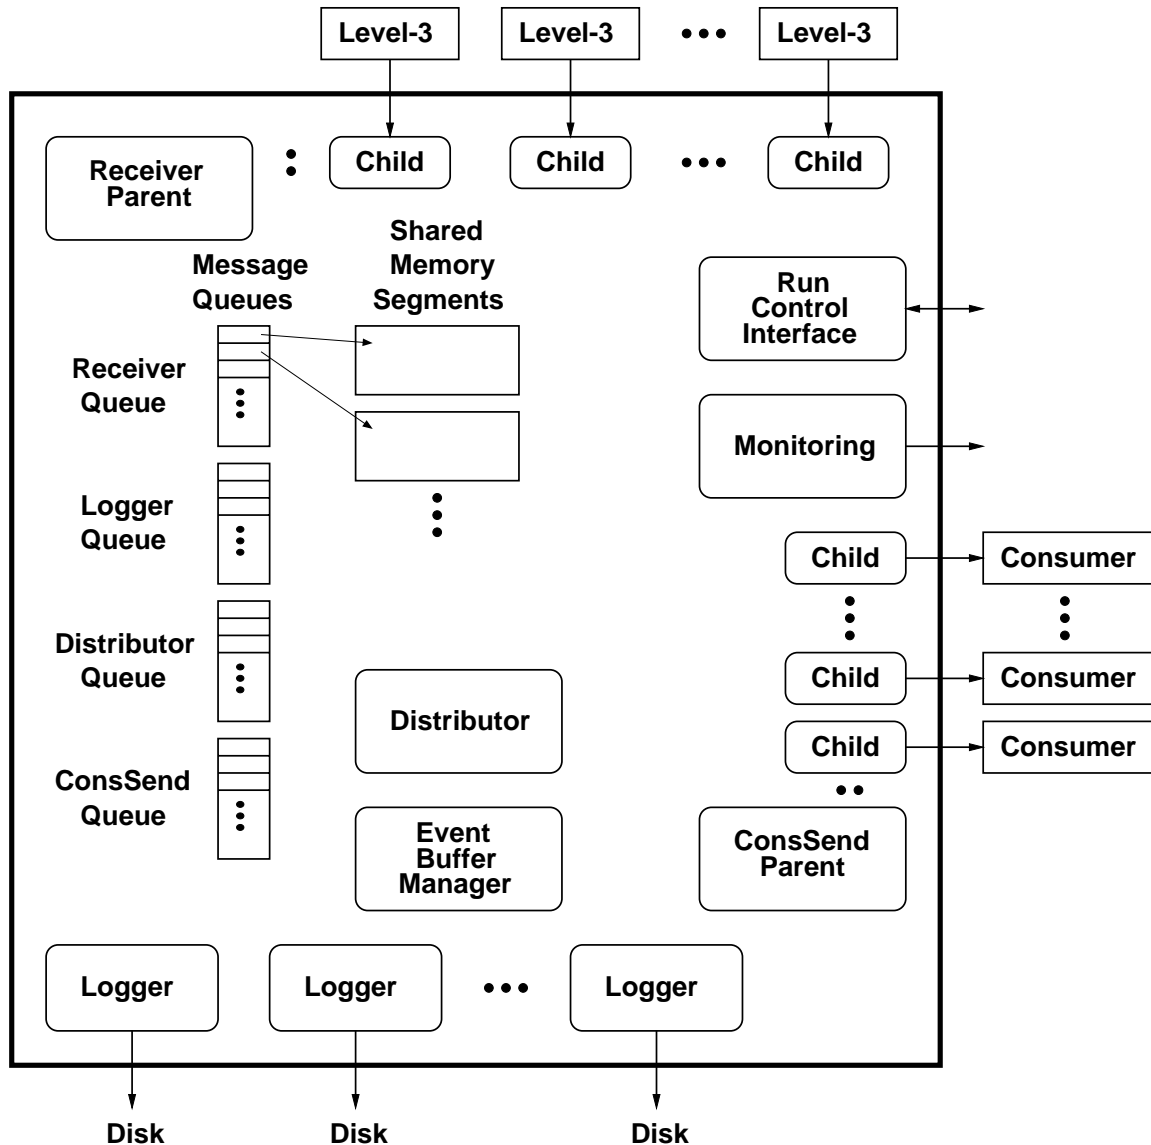
describing why the CSL was restarted.

CSL hardware overview



- b0dau32, an SGI 2200 Server (4 CPUs, large I/O bandwidth) located on third floor of B0
- ~3.5 TB of disk space on third floor of B0 (7 RAIDs)
- dual ported disks allow both fcdfsgi1 in FCC and b0dau32 to access disks

CSL software: overview



The CSL is a “server”. Possible “clients” include

- Level-3 output node processes
- software event builder processes
- 24 hour sender in partition 14
- consumers

Events are stored in shared memory buffers. Flow of events between processes inside CSL achieved by means of message queues.

CSL software: monitoring

The monitoring process collects CSL status information and sends it via a smartsockets message to the CSL monitoring display.

There are two kinds of CSL monitoring:

- the CSL display (snapshot of current CSL state)
- the CSL history plots (show rates, disk space, events logged, etc. versus time)

Using CSL display to check current CSL status:

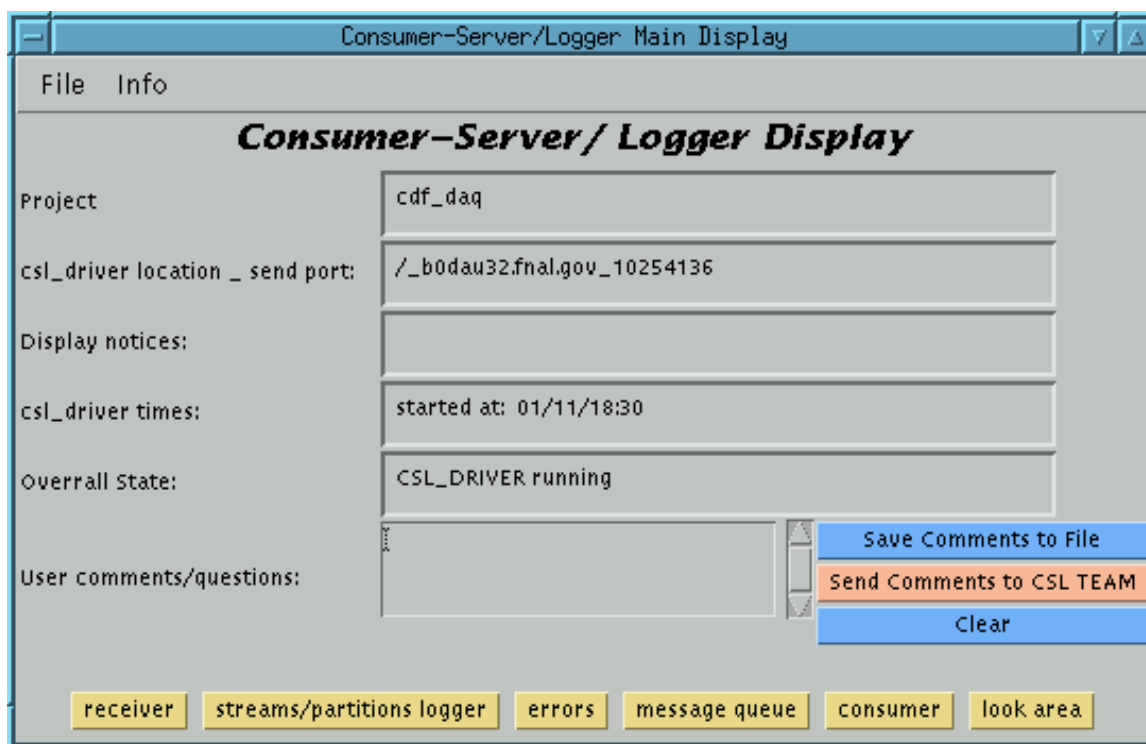
login to machine on online cluster

setup fer

daqmon

A GUI should appear on your screen. Click on CSL button.

Use project name cdfdaq (this should be the default). The CSL monitor main display window should appear.



CSL software: receiving

RECEIVER DISPLAY OF EVENTS COMING INTO CSL											
File		Info									
Events From			Statistics		Status						
Client Name	Partit.	Port #	Inst Rate	Event count	Event size	Activity rate	Unconn	Conn	Idle	Active	
b0dap60.fnal.gov	14	1435	0.235	949	0.297	0.0833					
b0l3u02c	0	4022	1.3121	242	0.3403	0.0833					
b0l3u01c	0	3023	1.3148	236	0.341	0.0667					
b0l3u07d	0	3523	1.2062	233	0.3297	0.0667					
b0l3u06c	0	4540	1.3542	207	0.334	0.05					
b0l3u04d	0	3452	1.4315	221	0.3367	0.05					
b0l3u08d	0	2361	1.2629	220	0.3452	0.0333					
b0l3u05c	0	4349	1.3569	222	0.3347	0.0833					
b0l3u03d	0	4097	1.2023	232	0.3474	0.0167					
Total (MB/s)		10.6759	Receivers ending in error			440					
Events Received from current clients		2762	Show average rate								
Events Received from all clients		0	Show current rate								

For each client sending events to the CSL, there is a receiver process.

The CSL monitoring GUI displays information for each receiver process: client node name, partition, number of events received, and rate.

CSL software: logging

The screenshot shows a window titled "rates for logger streams and partitions". It contains a table with columns for "Part 0" through "Part 15" and rows for "Run Number", "stream A" through "stream J", and "Part totals". The "Run Number" row shows values for each partition, with Part 0 being 136395 and others being 0. The "stream" rows show values for each stream, with stream A being 49, stream B being 2178, stream G being 491, and stream J being 532, while others are 0. The "Part totals" row shows a total of 3250 for Part 0 and 0 for all other partitions. At the bottom, there are buttons for "Show peak rate", "Show average rate", "Show current rate", "Show number of events", and "Show event rate".

	Part 0	Part 1	Part 2	Part 3	Part 4	Part 5	Part 6	Part 7	Part 8	Part 9	Part 10	Part 11	Part 12	Part 13	Part 14	Part 15
Run Number	136395	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream A	49	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream B	2178	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream C	0	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream D	0	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream E	0	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream F	0	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream G	491	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream H	0	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream I	0	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
stream J	532	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000	0	0.000000
Part totals	3250	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Total				3250												
Overall Logging counter																

This display is one way to check whether events are being written to disk.

For each partition and each stream this display shows

- the number of events written to disk OR
- the rate of writing events to disk in MB/sec

Modes of running: 4 stream (A,B,G,J), inclusive (only stream I), 8 stream (A,B,C,D,E,G,H,J)

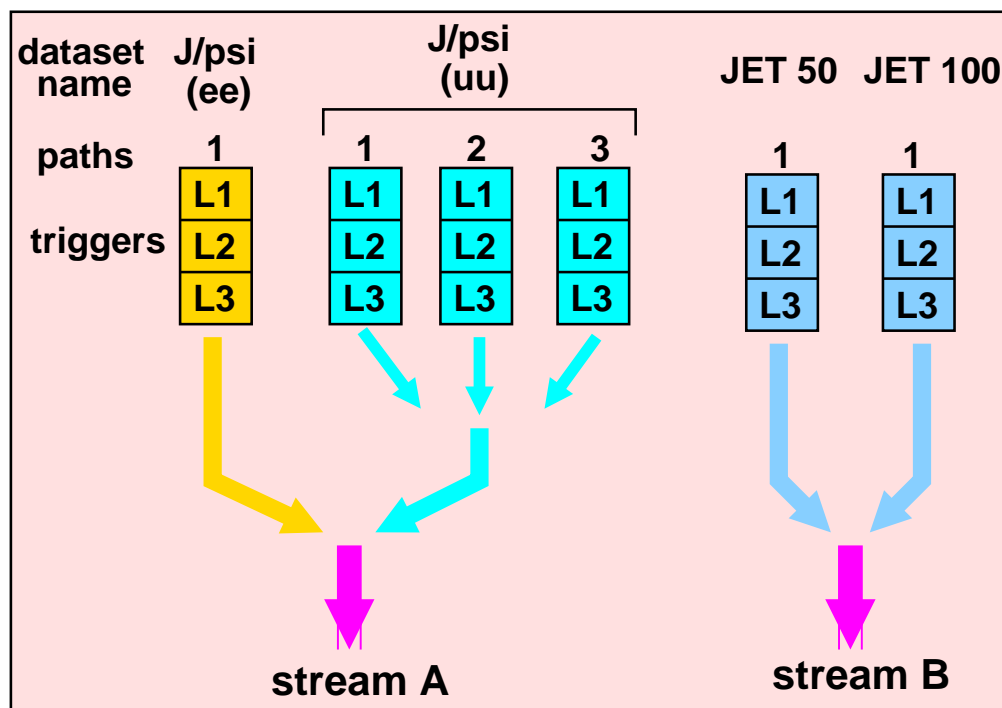
CSL writes events into 1 GB output files onto disks.

CSL disks are temporary holding space for events.

After closing, file sent to FCC where fcdfsi1 puts it onto tape. Should be enough disk space for an 8 hour buffer.

First 1 or 2 files of each run go to “look” area on fcdfsi2.

Paths, datasets, streams in Run 2



path: AND of Level-1, Level-2, Level-3 triggers.

dataset: OR of all paths defined for that dataset.

stream: collection of datasets

Events coming out of Level-3 are “streamed”: tagged as belonging to a particular stream.

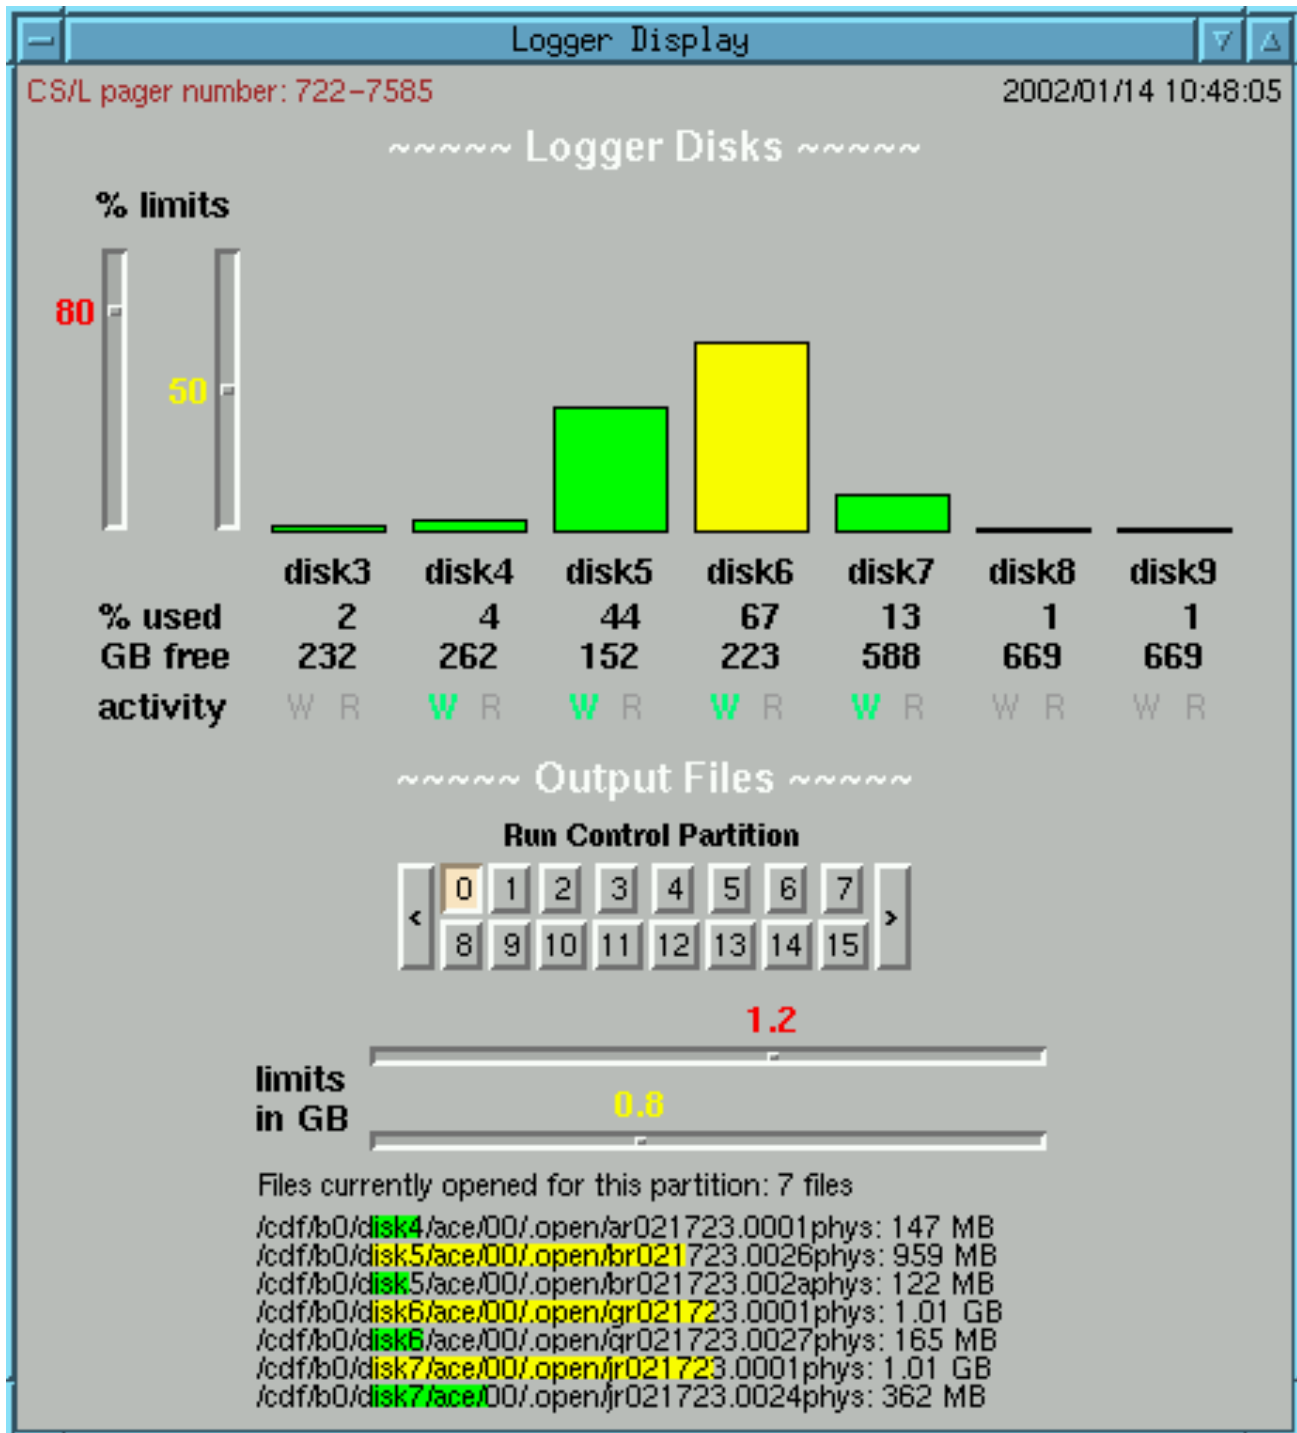
- CDF DAQ system can run in multi-partition mode
- each partition is independent of other partitions
- CSL writes events into separate files for different partitions
- CSL writes events corresponding to different streams within a partition into separate files

Run II: 8 streams, 50-100 datasets

CSL software: serving consumers


Consumer display										
Consumer		Statistics			Status					
Cons. Type	Cons. id / part. id	Host name or address	Evs Req.	Events Recvd	Inst Rate	Conn	Waiting	Active	Comple	
DAQMon	28214 / 0	b0dap69.fnal.gov	2502	2501	1.9820185					
Stage0	28212 / 0	b0dap69.fnal.gov	589	589	0.4119289					
SiliMon	25093 / 0	b0dap67.fnal.gov	70	70	0.1095379					
SVXMon	27223 / 0	b0dap66.fnal.gov	138	138	0.1098888					
ObjectMon	20062 / 0	b0dap56.fnal.gov	2005	2005	1.7567364					
BeamMon	27621 / 0	b0dap66.fnal.gov	2449	2448	1.9558252					
LumMon	18337 / 0	b0dap67.fnal.gov	2520	2520	1.992783					
EventDisplay	2429 / 0	b0dap50.fnal.gov	1	1	0					
TrigMon	31169 / 0	b0dap52.fnal.gov	591	591	0.4792719					
XMon	28354 / 0	b0dap70.fnal.gov	544	544	0.4473012					

CSL Monitoring: disks



Shows how full each CSL disk is, name and size of open files for each partition. This is one way to check that data is being written to disk. (Can also use “findfile runnumber” command on b0dau32.)

CSL Monitoring: message queues



There are 150 internal buffers in the CSL. Each buffer can store one event. The display shows four message queues. Each message on a queue points to a buffer where an event may be stored.

Receiver (left, blue) queue shows how many buffers are free to store new events from Level-3. If this is ALWAYS zero and "logger" queue has all the buffers then there is a problem. Contact CSL expert.

CSL history plots

Netscape: Consumer-Server/Logger Home Page

File Edit View Go Communicator Help

Back Forward Reload Home Search Netscape Print Security Shop Stop

Bookmarks Location: <http://www-b0.fnal.gov:8000/~csl/> What's Related

WebMail Contact People Yellow Pages Download Channels

[CSL home page](#)
[presentations](#)
[documentation](#)
[contact people](#)
[monitor plots](#)
[useful links](#)

Consumer-Server/Logger Monitoring Plots

CSL specific plots		
receiver, logger, consumer rates	24 hours	one week
number of events logged	24 hours and one week	
number of processes	24 hours	one week
number of messages in queues	24 hours	one week
average event size (partition 0-3)	24 hours	one week
average event size (partition 4-7)	24 hours	one week
CPU usage: receivers,loggers,consends	24 hours	one week
CPU usage: driver,distributor,bufman,monsend	24 hours	one week
look area	24 hours	one week
general b0dau32 plots		
disk space usage	24 hours	one week
load average	24 hours and one week	
free memory	24 hours and one week	
CPU usage: global	24 hours	one week

[List of files waiting to be copied to the lookarea](#)

[Accelerator Status](#)

Calibration CSL

A special version of the CSL software runs on b0dap60. Useful for some calibration runs which require

- guaranteed delivery of all events to the consumer OR
- a very large event size
(expected Run 2 event size is about 250 kB, official CSL can accept up to 3 MB, calibration CSL can accept up to 17 MB)

Calibration CSL does not log any data to disk.

Data File Catalog

CSL writes information into the Data File Catalog database for each output data file:

- file size
- total number of events in the file
- run number
- first event number in the file
- last event number in the file
- run section numbers

This must be done before the files are put onto tape.

Look Area on fcdsg12

To check status: on b0dau32 as user ace type “lookarea check”

Troubleshooting

Here are some reasons why you might think there is a problem with the CSL

- Level-3 is stuck in a “waiting to output events” state
- The CSL does not acknowledge a run control transition.
- There is an error message from the CSL in the Error Logger.
- You believe Level-3 or the software event builder is sending events to the CSL and one or more of the following is true
 - the consumers are not receiving any events
 - the events do not seem to be written to disk by the CSL
 - no files for the runs you are taking appear in the ”look” area on fcdfsgi2

Some things you should check before paging a CSL expert:

- Is the CSL receiving any events?
- If the CSL is not receiving events...
- Is the CSL sending events to consumers?
- Is the CSL writing events to disk?
- Did the CSL send an error message to the Error Logger?

CSL ACE web pages describes how to answer these questions.

AFTER going through the checklist, if you still think there is a problem with the CSL, then page a CSL expert.